

EFFICIENT 3D NLTE DUST RADIATIVE TRANSFER WITH SKIRT

MAARTEN BAES, JORIS VERSTAPPEN, ILSE DE LOOZE, JACOPO FRITZ, WAAD SAFTLY, EDGARDO VIDAL PÉREZ
 Sterrenkundig Observatorium, Universiteit Gent, Krijgslaan 281-S9, B-9000 Gent, Belgium

AND

MARKO STALEVSKI

Astronomical Observatory Belgrade, Volgina 7, P.O.Box 74 11060 Belgrade, Serbia
 Sterrenkundig Observatorium, Universiteit Gent, Krijgslaan 281-S9, B-9000 Gent, Belgium and
 Isaac Newton Institute, Yugoslavia branch

AND

SANDER VALCKE

Applied Maths NV, Keistraat 120, 9830 Sint-Martens-Latem, Belgium and
 Sterrenkundig Observatorium, Universiteit Gent, Krijgslaan 281-S9, B-9000 Gent, Belgium
Draft version August 26, 2011

ABSTRACT

We present an updated version of SKIRT, a 3D Monte Carlo radiative transfer code developed to simulate dusty galaxies. The main novel characteristics of the SKIRT code are the use of a stellar foam to generate random positions, an efficient combination of eternal forced scattering and continuous absorption, and a new library approach that links the radiative transfer code to the DustEM dust emission library. This approach enables a fast, accurate and self-consistent calculation of the dust emission of arbitrary mixtures of transiently heated dust grains and polycyclic aromatic hydrocarbons, even for full 3D models containing millions of dust cells. We have demonstrated the accuracy of the SKIRT code through a set of simulations based on the edge-on spiral galaxy UGC 4754. The models we ran were gradually refined from a smooth, 2D, LTE model to a fully 3D model that includes NLTE dust emission and a clumpy structure of the dusty ISM. We find that clumpy models absorb UV and optical radiation less efficiently than smooth models with the same amount of dust, and that the dust in clumpy models is on average both cooler and less luminous. Our simulations demonstrate that, given the appropriate use of optimization techniques, it is possible to efficiently and accurately run Monte Carlo radiative transfer simulations of arbitrary 3D structures of several million dust cells, including a full calculation of the NLTE emission by arbitrary dust mixtures.

Subject headings: radiative transfer – galaxies: ISM – dust, extinction – infrared: galaxies – galaxies: individual: UGC 4754

1. INTRODUCTION

The effects of absorption and scattering by interstellar dust grains on the structural parameters of galaxies has been a long-standing and controversial issue. The only way to tackle this problem is to properly solve the continuum radiative transfer equation, taking into account realistic geometries and the physical processes of absorption and multiple anisotropic scattering. Over the years, many different and complementary approaches have been developed to tackle the continuum radiative transfer problem in simple geometries such as spherical or plane-parallel symmetry. While one-dimensional radiative transfer calculations have been crucial to isolate and demonstrate the often counter-intuitive aspects of important parameters such as star-dust geometry, dust scattering properties and clumping (Bruzual A. et al. 1988; Disney et al. 1989; Witt et al. 1992; di Bartolomeo et al. 1995; Baes & Dejonghe 2001a; Inoue 2005), we need more sophisticated radiative transfer models to model complicated systems such as disc galaxies in detail. Thanks to new techniques and ever increasing computing power, the construction of 2D and 3D realistic radiative transfer models is now possible (e.g. Xilouris et al. 1999; Popescu et al. 2000; Gordon et al. 2001; Steinacker et al. 2003;

Semionov & Vansevičius 2005; Pinte et al. 2006; Jonsson 2006; Bianchi 2008).

A complementary and powerful way to study the content of galaxies is to use the direct emission of dust at long wavelengths. Large dust grains will typically reach a state of local thermal equilibrium (LTE) in the local interstellar radiation field (ISRF) and re-radiate the absorbed UV/optical radiation at far-infrared (FIR) and submm wavelengths. Thanks to the spectacular advances in instrumentation in the FIR/submm wavelength region, we have seen a significant improvement in the amount of FIR/submm data on both nearby and distant galaxies. In particular, the launch of the Herschel Space Observatory with the sensitive PACS and SPIRE instruments has enabled both the detailed study of nearby, resolved galaxies (Bendo et al. 2010; Galametz et al. 2010; Pohlen et al. 2010; Roussel et al. 2010; Smith et al. 2010) and the detection of thousands of distant galaxies (Clements et al. 2010; Oliver et al. 2010; Rigby et al. 2011). Whereas large grains typically emit as a modified blackbody at an equilibrium temperature of 15-30 K and hence dominate the FIR/submm emission of galaxies, small grains and PAH molecules are transiently heated by the absorption of single UV photons to much higher temperatures. The NLTE emission from very small grains and PAHs dominates the emission of

galaxies at mid-infrared wavelengths. The ISO and particularly the Spitzer mission have been instrumental in uncovering the mid-infrared emission of nearby galaxies (Helou et al. 2000; Smith et al. 2007; Draine et al. 2007).

Different approaches have been developed to calculate the NLTE emission spectrum due to very small grains and PAHs (Dwek 1986; Désert et al. 1986; Guhathakurta & Draine 1989; Siebenmorgen et al. 1992; Draine & Li 2001; Compiègne et al. 2011), but the integration of NLTE emission into radiative transfer codes has proven to be a challenging task. The main reason is that the computational effort necessary to calculate the temperature distribution of the different dust grains is substantial. In the general case, the calculation of the dust emissivity in a single dust cell requires the solution of a large matrix equation for each single dust population, with the size of the matrix determined by the number of temperature or enthalpy bins. In the so-called thermal continuous cooling approximation (Draine & Li 2001), this matrix equation can be solved recursively, but still the calculation of the emission spectrum remains a significant computational challenge. Indeed, since the temperature distribution of dust grains depends strongly on both the size of the grains and the strength and hardness of the ISRF, a large number of temperature or enthalpy bins is necessary to sample the temperature distribution correctly. Moreover, because of this strong dependence on grain size and ISRF, the choice of the temperature bins is hard to fix a priori and an iterative procedure is to be preferred.

In spite of the high numerical cost, NLTE dust emission has been built into several radiative transfer codes, using various approximations and/or assumptions. The most simple approach is the one followed by e.g. Wood et al. (2008) and Jonsson et al. (2010), who use a set of predefined NLTE dust emissivities with the simplifying assumption that the emissivity is a function only of strength and not of the spectral shape of the exciting ISRF. A pioneering code in which NLTE dust emission was included in a self-consistent way was the 2D ray-tracing code by Siebenmorgen et al. (1992). The number of temperature distribution calculations are minimized by the assumptions that grains with a size larger than about 80 Å are in thermal equilibrium, and by the use of a pre-fixed time dependence of the cooling of PAH grains. A similar approach was adopted in the 3D Monte Carlo radiative transfer code DIRTY (Gordon et al. 2001; Misselt et al. 2001). The TRADING code by Bianchi (2008) uses a different approach: this code uses a fixed (and limited) grid of temperature bins for all ISRFs and grain sizes, which allows to precompute and tabulate a significant fraction of the quantities necessary for the calculation of the temperature distribution. Yet a different approach is the work by Juvela & Padoan (2003): driven by the observation that the spectrum of the local ISRF is very similar in many places in a dusty medium, they considered the idea of a dynamic library of dust emission spectra. The idea is that the intensity of the ISRF at a very limited number of reference wavelengths (they typically used only two) suffices to make a reliable estimate of the total ISRF and hence of the dust emission spectrum.

In this paper we present an updated version of the SKIRT Monte Carlo radiative transfer code. This code,

of which the name is an acronym to Stellar Kinematics Including Radiative Transfer, was initially developed to study the effect of dust absorption and scattering on the observed kinematics of dusty galaxies (Baes & Dejonghe 2001b, 2002; Baes et al. 2003). In a second stage, the SKIRT code was extended with a module to self-consistently calculate the dust emission spectrum under the assumption of local thermal equilibrium (Baes et al. 2005a,b). This LTE version of SKIRT has been used to model the dust extinction and emission of various types of galaxies (Gomez et al. 2010; Baes et al. 2010; De Looze et al. 2010), as well as circumstellar discs (Vidal & Baes 2007) and clumpy tori around active galactic nuclei (Stalevski et al. 2011; Popović et al. 2011). In this present paper we present a strongly extended version of the SKIRT code that can perform efficient 3D radiative transfer calculations including a self-consistent calculation of the dust temperature distribution and the associated FIR/submm emission with a full incorporation of the emission of transiently heated grains and PAH molecules.

In Section 2 we present the general characteristics of the SKIRT code, whereas we highlight a number of particular aspects in Section 3 and some implementation details in Section 4. In Section 5 we describe a number of tests and applications, and Section 6 we present our conclusions.

2. THE SKIRT MONTE CARLO RADIATIVE TRANSFER CODE

2.1. Monte Carlo radiative transfer

SKIRT is a 3D continuum radiative transfer code based on the Monte Carlo algorithm. The key principle in Monte Carlo radiative transfer simulations is that the radiation field is treated as a flow of a finite number of photon packages. A simulation consists of consecutively following the individual path of each single photon package through the dusty medium. The journey or lifetime of a single photon package can be thought of as a loop: at each moment in the simulation, a photon package is characterized by a number of properties, which are generally updated when the photon package moves to a different stage on its trajectory. The trajectory of the photon package is governed by various events such as emission, absorption and scattering events. Each of these events is determined statistically by random numbers, generated from the appropriate probability distribution $p(x) dx$. Typically, a photon is emitted by a star, undergoes a number of scattering events and its journey ultimately ends when it is either absorbed by the dust or it leaves the system. A Monte Carlo simulation repeats this same loop for every single one of the photon packages and analyzes the results afterwards.

The mathematical details and practical implementation of Monte Carlo radiative transfer have both been described in detail by various authors (e.g. Cashwell & Everett 1959; Mattila 1970; Witt 1977; Fischer et al. 1994; Bianchi et al. 1996; Gordon et al. 2001; Niccolini et al. 2003; Wolf 2003; Stamatellos & Whitworth 2003; Juvela 2005; Jonsson 2006; Bianchi 2008) and will not be repeated here in full detail. Our overall approach is comparable to the DIRTY (Gordon et al. 2001; Misselt et al. 2001) and TRADING (Bianchi 2008) radiative transfer codes and we refer the interested reader to these papers

for more details. We will only give a compact description of the general characteristics of the SKIRT Monte Carlo code and not describe all the details. Instead, we will focus our attention to those aspects of the SKIRT code that are novel and/or different compared to the other codes.

2.2. General overview of a SKIRT simulation

Each SKIRT simulation consists of four phases: the initialization phase, the stellar emission phase, the dust emission phase, and the clean-up phase.

2.2.1. The initialization phase

The initialization phase consists of adopting the correct unit system, setting up the random number generator, computing the optical properties of the various dust populations, constructing the dust grid, setting up the stellar geometry and setting up the instruments of the various observers. Once this initialization is finished, the actual simulation can start.

2.2.2. The stellar emission phase

In the stellar emission phase, we consider the transfer of the primary source of radiation (usually stellar sources, but it can also include an accretion disc or nebular line emission) through the dusty medium. The stellar emission phase consists of a set of parallel loops, each of them corresponding to a single wavelength. At every single wavelength, the total luminosity of the stellar system is divided into a very large number (typically 10^5 to 10^7) of monochromatic photon packages, which are launched consecutively through the dusty medium in random propagation directions.

Once a photon package is launched into the dusty medium (either after an emission event or following a scattering event), it can be absorbed by a dust grain, it can be scattered by a dust grain, or it can travel through the system without any interaction. In a naive Monte Carlo routine, these three possibilities are possible and it is randomly determined which of the three will happen. This is generally an inefficient procedure, though, which leads to poor signal-to-noise both in the absorption rates in the different cells and in the scattered light images. To overcome these problems, we have set up a combination of continuous absorption and eternal forced scattering (see Section 3.3 for details). The result is that, contrary to most Monte Carlo codes where the life cycle of a photon package ends when it either leaves the system or is absorbed, the photon packages in SKIRT can never leave the system. The life cycle of a photon package ends when the package contains virtually no more luminosity (typically we use the criterion that it must have lost 99.99% of its original luminosity). Whenever this happens, a different stellar photon package is launched until also this one finishes its life cycle. This loop is repeated for all stellar photon packages at a given wavelength, and subsequently for all wavelengths (in a multi-core system, each core can handle the loop corresponding to a different wavelength at the same time).

2.2.3. The dust emission phase

After the stellar emission phase, the code moves to the dust emission phase. This phase is roughly similar to

the stellar phase, except that the sources that emit the radiation are now not the primary, stellar sources but the dust cells. From the stellar emission phase we know the total amount of absorbed radiation at each wavelength in each cell of the dust domain. From this absorption rate we can calculate the mean intensity of the ISRF in each cell, which allows the calculation of the dust emissivity, depending on the physical processes the SKIRT user is interested in (see Section 3.5). With the sources (the dust cells) and their emissivity determined, the simulation now enters a loop that is very similar to the one in the stellar emission phase. At each individual wavelength, a huge number of photon packages is generated which are launched and followed consecutively through the dusty medium. Care is taken that all regions of the dusty medium, including those cells where only a small amount of luminosity has been absorbed, are well-sampled.

The dust-emitted photon packages in turn increase the absorption rate in the dust cells where they pass through. This results in an increase of the mean intensity of the ISRF. The result is that at the end of the dust emission phase, the absorption rates used to calculate the dust emissivity in each cell do not correspond to the mean intensity of the ISRF. This naturally leads to an iterative procedure, in which the absorption rate, the mean intensity and the dust emissivity are updated until convergence is achieved (Misselt et al. 2001; Bianchi 2008). We hence repeat the dust emission phase of the code several times. We typically require a 1% level accuracy in the dust bolometric absorption rate of each cell as a stopping criterion for the iteration. It is typically reached in only a few iterations; for all the simulations we have done so far, less than five iterations have been necessary.

2.2.4. The clean-up phase

The last phase of the Monte Carlo simulation starts when the last of the photon packages emitted by the dust component has lost 99.99% of its initial luminosity. It simply consists of calibrating and reading out the instruments (all output is written to FITS files) and other useful information, such as 3D absorption rate maps and dust temperature distributions.

3. PARTICULAR ASPECTS OF THE SKIRT CODE

3.1. Setup of the dust grid

A critical aspect in Monte Carlo radiative transfer simulations is the choice of the dust grid. The dust grid consists of tiny cells, each of which have a number of characteristics that fully describe the physical properties of the dust at the location of the cell. The choice of the grid has a significant impact on both the run time and the memory requirement of the simulation. Indeed, each photon package typically requires several integrations through the dust (i.e. the determination of the optical depth along the path and the conversion of a given optical depth to a physical path length), and the calculation time of a single optical depth typically scales with the number of grid cells crossed. Different kinds of dust grids can be applied in the SKIRT code. The most general grid is a 3D cartesian grid in which each dust cell is a rectangular cuboid. For simulations with a spherical or axial symmetry, we also have 1D spherical

and 2D cylindrical grids (the elementary dust cells being shells or tori respectively). The distribution of the grid points (in 1D spherical, 2D cylindrical or 3D cartesian grids) can be chosen arbitrarily; linear, logarithmic, or power-law cell distributions have been pre-programmed, but any user-supplied grid cell distribution is possible.

The main goal of the dust grid is to discretize the dust density. We assume that the density of each dust component is uniform within each individual cell. In principle, the density does not need to be constant within each dust cell (see e.g. [Nicolini et al. 2003](#)). In the first versions of SKIRT, we have experimented with a more sophisticated kind of dust grid, where the density of the dust within each cell is not uniform but determined by trilinear interpolation of the values of the density on the eight border points of the cell (in case of a cartesian grid with cubic cells). In this case, the computation of optical depths in the dusty medium take more computation time, but the accuracy is increased such that a grid with less cells and less photon packages are needed per simulation. For models in which only absorption and scattering are taken into account, we found that this kind of dust grid is computationally more efficient than a dust grid with uniform density, in particular when the system harbors a large dynamical range of dust densities ([Baes et al. 2003](#)). However, for radiative transfer simulations in which the thermal emission of the dust is taken into account, each dust cell needs to contain an absorption rate counter, which collects the absorbed luminosity at every wavelength. The size of the dust cells is hence the typical resolution of the simulation, and the advantage of the interpolated grid (where the dust grid can be coarser because the density is resolved within each cell) largely disappears. SKIRT therefore only uses dust grids with a uniform density in each cell.

3.2. Sampling the stellar density

The first step in the life cycle of each stellar photon is the random generation of the location where it is emitted. This means that we have to generate random positions from the three-dimensional probability distribution

$$p(\mathbf{x}) d\mathbf{x} = \frac{\nu_\lambda(\mathbf{x}) d\mathbf{x}}{\int \nu_\lambda(\mathbf{x}') d\mathbf{x}'} \quad (1)$$

where $\nu_\lambda(\mathbf{x})$ is the luminosity density at the wavelength λ of the photon package. As SKIRT is primarily focused towards modeling galaxies, we have done efforts to optimize the generation of random positions from arbitrary 3D probability functions. The SKIRT code contains a library with common geometries for which the generation of a random position vector can be performed analytically. These include spherical [Plummer \(1911\)](#), [Jaffe \(1983\)](#) or [Hernquist \(1990\)](#) models, or axisymmetric power-law, exponential, sech or isothermal disc models. For other frequently used luminosity density profiles, e.g. flattened [de Vaucouleurs \(1948\)](#) or [Sérsic \(1968\)](#) models, or more general density profiles that cannot be described by an analytical function, such a direct analytical inversion is not possible. Two complementary approaches have been included in the SKIRT code to deal with generating random positions from such density profiles.

3.2.1. Multi-gaussian expansion technique




FIG. 1.— Observed (top) and simulated (bottom) V-band images of the Sombrero Galaxy (M104). The observed image is taken from the Spitzer Infrared Nearby Galaxies Survey (SINGS, [Kennicutt et al. 2003](#)) ancillary data website. In the SKIRT simulation, the stellar density profile is based on a multi-gaussian expansion (MGE) with 15 components, as presented by [Emsellem \(1995\)](#). For a full-scale panchromatic radiative transfer modeling of stars and dust in M104, see [De Looze et al. \(2011\)](#).

The first technique is to expand the density profile into a set of subcomponents. SKIRT contains a routine to perform a multi-gaussian expansion (MGE) of surface brightness distributions. An MGE expansion basically expands any surface brightness distribution as a finite sum of two-dimensional gaussian components ([Monnet et al. 1992](#); [Emsellem et al. 1994a](#); [Cappellari 2002](#)). The MGE method has proven to be a very powerful tool for image analysis: even with a relatively modest set of gaussian components, $N \sim 10$, even complex geometries can be reproduced accurately ([Emsellem et al. 1994b, 1999](#); [Cappellari 2002](#); [Cappellari et al. 2006](#); [van den Bosch et al. 2008](#)). One of the reasons why a MGE expansion is very useful for SKIRT is that this approach enables a straightforward determination of the 3D spatial distribution: if the Euler angles of the line of sight are known, the de-projection of a 2D gaussian distribution on the plane of the sky is a 3D gaussian distribution and the conversion formulae are completely analytical ([Monnet et al. 1992](#)). Generating random positions from a sum of 3D gaussian probability distributions is straightforward. An example of this approach can be seen in Figure 1, where we present a radiative transfer model for the Sombrero Galaxy based on the MGE expansion of its surface brightness distribution presented by [Emsellem \(1995\)](#).

3.2.2. The stellar foam

The second approach consists of sampling random positions directly from the stellar density distribution using the so-called stellar foam. The stellar foam is a SKIRT structure based on the Foam library developed by [Jadach \(2003\)](#) and [Jadach & Skrzypek \(2007\)](#). Foam is a self-adapting cellular Monte Carlo tool aimed at Monte Carlo

integration of multi-dimensional functions, including integrands with an arbitrary pattern of singularities. It achieves a high efficiency thanks to an intelligent division of the integration space into small simplicial or hyper-rectangular form, which are created in a self-adaptive way by binary splitting. It has originally been developed for use in high-energy physics (e.g. [Jadach & Skrzypek 2006](#); [Andonov et al. 2010](#); [Haas & Makarenko 2001](#)), but it can also be adopted as a general-purpose Monte Carlo event generator.

We use an adapted version of the Foam library for the generation of random positions from an arbitrary probability density $p(\mathbf{x}) d\mathbf{x}$. One problem is that the probability density $p(\mathbf{x}) d\mathbf{x}$ from which we want to generate random positions is typically defined on the entire 3D space, whereas Foam requires a probability density on the N -dimensional unit hypercube $[0, 1]^N$. We achieve this through a coordinate transformation from the usual coordinates $\mathbf{x} = (x, y, z)$ to new coordinates $\bar{\mathbf{x}} = (\bar{x}, \bar{y}, \bar{z})$ such that we map each infinite interval $[-\infty, \infty]$ onto the unit interval $[0, 1]$. The probability density $p(\mathbf{x}) d\mathbf{x}$, normalized such that

$$\int_{-\infty}^{\infty} dx \int_{-\infty}^{\infty} dy \int_{-\infty}^{\infty} p(x, y, z) dz = 1 \quad (2)$$

is transformed to a new probability density

$$\bar{p}(\bar{\mathbf{x}}) d\bar{\mathbf{x}} = p\left(x(\bar{x}), y(\bar{y}), z(\bar{z})\right) \left| \frac{dx}{d\bar{x}} \frac{dy}{d\bar{y}} \frac{dz}{d\bar{z}} \right| d\bar{\mathbf{x}} \quad (3)$$

This new density will be defined and normalized to one on the unit cube,

$$\int_0^1 d\bar{x} \int_0^1 d\bar{y} \int_0^1 \bar{p}(\bar{x}, \bar{y}, \bar{z}) d\bar{z} = 1 \quad (4)$$

There are many possible transformations we can apply to achieve this. We have chosen the transformation

$$\bar{x} = \frac{1}{\pi} \arctan\left(\frac{a}{x}\right) \quad (5a)$$

$$\bar{y} = \frac{1}{\pi} \arctan\left(\frac{b}{y}\right) \quad (5b)$$

$$\bar{z} = \frac{1}{\pi} \arctan\left(\frac{c}{z}\right) \quad (5c)$$

with a , b and c three scale parameters, which we can adapt for the specific probability density we are considering. The inverse transformation is

$$x = \frac{a}{\tan(\bar{x}\pi)} \quad (6a)$$

$$y = \frac{b}{\tan(\bar{y}\pi)} \quad (6b)$$

$$z = \frac{c}{\tan(\bar{z}\pi)} \quad (6c)$$

and the Jacobian reads

$$\left| \frac{\partial(x, y, z)}{\partial(\bar{x}, \bar{y}, \bar{z})} \right| = \left| \frac{dx}{d\bar{x}} \frac{dy}{d\bar{y}} \frac{dz}{d\bar{z}} \right| = \frac{abc \pi^3}{\sin^2(\bar{x}\pi) \sin^2(\bar{y}\pi) \sin^2(\bar{z}\pi)} \quad (7)$$

Summarizing, if we want to generate random positions \mathbf{x} from an arbitrary probability density $p(\mathbf{x}) d\mathbf{x}$, we first determine representative scale lengths a , b and c along the three dimensions, and subsequently calculate the corresponding probability density $\bar{p}(\bar{\mathbf{x}}) d\bar{\mathbf{x}}$ using equations (3),

(6) and (7). The foam generator is applied to this new probability function to generate the random points $\bar{\mathbf{x}}$, which are converted to the desired positions \mathbf{x} through the formulae (6). The construction of the stellar foam takes only a few seconds in two dimensions up to about one minute in three dimensions.

3.3. Eternal forced scattering and continuous absorption

Once a photon package has been generated at a random location (sampled from the stellar density) and it has been given a random propagation direction (sampled randomly from the unit sphere), it is ready to start its journey through the dusty medium. It has three possible fates: it can be absorbed by a dust grain at a certain position along its path, it can be scattered by a dust grain at a certain position on its path, or it can travel along its path through the system without any interaction. The probability for each of these three options are determined by the scattering albedo ϖ_λ and the optical depth $\tau_{\lambda, \text{path}}$ along the path, defined as

$$\tau_{\lambda, \text{path}} = \int_0^\infty \kappa_\lambda \rho(s) ds \quad (8)$$

where κ_λ is the extinction coefficient, ρ is the dust density and the integral covers the entire path of the photon package through the dusty medium. The most straightforward way to model these different physical processes in a Monte Carlo radiative transfer code is to generate randomly which of these three processes will take place. In case the photon package leaves the system, its lifetime is terminated and a different package is launched. In the case of a scattering event, the position of the scattering event is determined by choosing a random path length from the appropriate probability distribution and a scattering event is simulated. Finally, in case it is absorbed, the position of the absorption is determined in a similar way, the luminosity of the photon package is stored in a local absorption rate counter attached to the dust cell where the absorption event took place, and the photon package's life is terminated. These local absorption rates are used in a later stage of the simulation to estimate the local mean intensity of the ISRF, necessary to calculate the thermal dust emission.

This traditional method has two significant drawbacks: along paths where the optical depth is modest, many photon packages will escape from the system without interactions, which will result in bad statistics of the scattered intensity and the absorbed luminosity. Even if the photon packages do interact, most interactions will take place on those sections of the path where the density is largest. Many absorption events are necessary in each cell to guarantee a high-quality estimate of the local absorption rate and the corresponding mean intensity. In dust cells with a low density (such that only few absorptions take place) and at wavelengths where the absorption rate is low, this usually requires large numbers of photon packages.

These two problems can be minimized using the efficient combination of two clever Monte Carlo techniques: forced scattering and continuous absorption. Continuous absorption (or Lucy estimation) is a technique to estimate the mean intensity of the ISRF throughout the dusty medium ([Lucy 1999](#); [Niccolini et al. 2003](#)). The continuous absorption technique is designed to solve the

problem of poor statistics in the absorption rate in low-density regions by spreading the absorption over all cells along the photon package's path instead of concentrating it in one single cell. A different but equivalent way to see it is that the mean intensity in each cell is estimated using the sum of the path lengths covered by all photon packages that traverse that particular cell. Forced scattering is an old technique that was already implemented in the first Monte Carlo radiative transfer codes (Cashwell & Everett 1959; Mattila 1970; Witt 1977). When applying forced scattering, photon packages are forced to interact with a dust grain before they leave the system. This incorrect behavior is corrected for by decreasing the luminosity of the photon package.

In most radiative transfer codes, forced scattering is considered only after an emission event and subsequent scattering events are unforced. In the SKIRT code we always consider forced scattering, such that we have eternal forced scattering. The combination of eternal forced scattering and continuous absorption results in a very efficient Monte Carlo routine. Instead of determining randomly whether a photon package with luminosity L_λ will escape, will be absorbed or will be scattered, we split it into $n + 2$ child photon packages (with n the number of dust cells along the path): one child photon package that will leave the system without interaction, one child photon package that will be scattered by a dust grain somewhere along the path, and n children that will be absorbed, one in each of the n cells along the path. The luminosity of each of these children is easy to calculate: we find

$$L_{\lambda, \text{abs}_1} = L_\lambda (1 - \varpi_\lambda) (1 - e^{-\tau_{\lambda,1}}) \quad (9)$$

$$L_{\lambda, \text{abs}_2} = L_\lambda (1 - \varpi_\lambda) (e^{-\tau_{\lambda,1}} - e^{-\tau_{\lambda,2}}) \quad (10)$$

⋮

$$L_{\lambda, \text{abs}_{n-1}} = L_\lambda (1 - \varpi_\lambda) (e^{-\tau_{\lambda,n-2}} - e^{-\tau_{\lambda,n-1}}) \quad (11)$$

$$L_{\lambda, \text{abs}_n} = L_\lambda (1 - \varpi_\lambda) (e^{-\tau_{\lambda,n-1}} - e^{-\tau_{\lambda, \text{path}}}) \quad (12)$$

$$L_{\lambda, \text{sca}} = L_\lambda (1 - \varpi_\lambda) (1 - e^{-\tau_{\lambda, \text{path}}}) \quad (13)$$

$$L_{\lambda, \text{esc}} = L_\lambda e^{-\tau_{\lambda, \text{path}}} \quad (14)$$

with $\tau_{\lambda,j}$ the optical depth to the point on the path where the photon package would leave the j 'th dust cell. Obviously, we have

$$L_{\lambda, \text{esc}} + L_{\lambda, \text{sca}} + \sum_{j=1}^m L_{\lambda, \text{abs}_j} = 1 \quad (15)$$

In our Monte Carlo simulation, we now consider each of these $n + 2$ child photon packages. Each one of the n children with luminosity $L_{\lambda, \text{abs}_j}$ is absorbed in the j 'th cell along the path, which means that its luminosity is added to the absorption rate counter attached to this dust cell. The child photon package with luminosity $L_{\lambda, \text{esc}}$ escapes from the system, which implies that we don't have to take this one into account anymore. Finally, the child photon package with luminosity $L_{\lambda, \text{sca}}$ is scattered somewhere along the path. This is basically the only photon package that we need to follow up. We still have to randomly determine the location of the scattering event along the path. This is achieved by selecting a random optical depth τ_λ from an exponential distribution over the finite

range $0 < \tau_\lambda < \tau_{\lambda, \text{path}}$, i.e.

$$p(\tau_\lambda) d\tau_\lambda = \frac{e^{-\tau_\lambda} d\tau_\lambda}{1 - e^{-\tau_{\lambda, \text{path}}}} \quad (16)$$

and translate this randomly generated τ_λ to a physical path length by solving the equation

$$\tau_\lambda = \int_0^s \kappa_\lambda \rho(s') ds' \quad (17)$$

for s . Once this path length has been determined, we can determine the position of the scattering event. If we then also determine a new propagation direction, determined randomly by generating a random direction from the scattering phase function, we are back at the starting point. This child now becomes the parent, it can be split into children and we can repeat the same loop all over again.

Summarizing, the net result of the combination of continuous absorption and eternal forced scattering is that after each emission/scattering event, we distribute a fraction of the photon package's luminosity among all the cells along the path, and we continue the Monte Carlo loop with a less luminous photon package that is always scattered at some point along the path. Hence, contrary to most Monte Carlo codes where the life cycle of a photon package ends when it either leaves the system or is absorbed, the photon packages in SKIRT can never leave the system. The life cycle of a photon package ends when the package contains virtually no more luminosity. Typically we use the criterion that it must have lost 99.99% of its original luminosity, which immediately is the minimum level of absolute energy conservation of the simulation.

3.4. Peeling off and smart detectors

The goal of a Monte Carlo radiative transfer simulation is to simulate observable properties of a dusty system, i.e. images and spectral energy distributions. SKIRT uses the technique of peel-off photon packages to create an arbitrary number of images/SEDs at different observing positions. Peeling off is a Monte Carlo technique designed to create high signal-to-noise images (in particular scattered light images), adopted for the first time by Yusef-Zadeh et al. (1984) and included in almost all state-of-the-art Monte Carlo radiative transfer codes. After every emission or scattering event, the code calculates which fraction of the luminosity contained in the photon package would arrive at the observers' locations and at which point on the plane of the sky, if the photon package would be emitted or scattered in the appropriate propagation direction. Repeating this for every photon package at every emission or scattering event implies that the maximum available information is obtained for a fixed set of photon packages and hence strongly increases the signal-to-noise compared to the more simple Monte Carlo codes where only photon packages that leave the system are recorded.

Each SKIRT detector is basically an integral field detector, i.e. a data cube with two spatial and one wavelength dimension. In most Monte Carlo radiative transfer codes the simulated detectors are natural, idealized representations of actual CCD detectors (or a series of them at each wavelength). They basically consist of a

two- or three-dimensional array of pixels, which act as a reservoir for the incoming photon packages. When a photon package leaves the system and arrives at the location of the observer, the correct pixel is determined and the luminosity of the photon package is added to the luminosity at that pixel. At the end of the simulation, the detector is read out pixel by pixel and the surface brightness distribution is constructed. While this approach seems the most natural way to simulate the detection of photon packages in a Monte Carlo simulation, it might not be the most efficient. We must be aware that, although we are simulating a real detection as closely as possible, we have more information at our disposal than real observers. The maximum information that a real observer can obtain (in the theoretical limit of perfect noise-free observations and instruments) when imaging with a CCD detector is the number of photon packages that arrive in each of his pixels. As numerical simulators, we have at our disposal the full information on the precise location of the impact of each photon package on the detector. In order to use this information, we have considered the concept of smart detectors, which take full advantage of the exact location of the impact of the incoming photon packages (Baes 2008). The principle of these smart detectors is based on the estimate of the density distribution in smoothed particle hydrodynamics simulations. While preserving the same effective resolution and having virtually no computational overhead, smart detectors realize a noise reduction of about 10 percent.

3.5. The dust emissivity

A crucial aspect of the SKIRT code is the calculation of the dust emissivity. From the stellar emission phase we know the total amount of absorbed radiation $L_{\lambda,m}^{\text{abs}}$ at each wavelength in each cell of the dust domain. From this absorption rate we can calculate the mean intensity of the ISRF $J_{\lambda,m}$ in each cell using

$$J_{\lambda,m} = \frac{L_{\lambda,m}^{\text{abs}}}{4\pi \kappa_{\lambda}^{\text{abs}} \rho_m V_m} \quad (18)$$

where $\kappa_{\lambda}^{\text{abs}}$ is the dust absorption coefficient, ρ_m is the dust density and V_m is the volume in cell number m respectively. Knowledge of the mean intensity and the dust properties in each cell allows the dust emissivity $j_{\lambda,m}^{\text{d}}$ to be determined. SKIRT contains three different modules for the calculation of the dust emissivity, depending on the physical processes that are taken into account.

3.5.1. Three different options for the dust emissivity

The simplest option is to consider the dust grains as a single species that is in local thermal equilibrium (LTE) with the local ISRF. In this case, the dust emits as a modified blackbody radiator,

$$j_{\lambda,m}^{\text{d}} = 4\pi \rho_m V_m \kappa_{\lambda}^{\text{abs}} B_{\lambda}(T) \quad (19)$$

where the dust equilibrium temperature T_m of the m 'th dust cell is determined by the condition of thermal and radiative equilibrium,

$$\int_0^{\infty} \kappa_{\lambda}^{\text{abs}} J_{\lambda,m} d\lambda = \int_0^{\infty} \kappa_{\lambda}^{\text{abs}} B_{\lambda}(T_m) d\lambda \quad (20)$$

The second, somewhat more realistic option is to still consider LTE for the dust grains, but taking into account that each species and size of dust grain reaches its

own equilibrium temperature. The SKIRT code allows to consider dust mixtures with an arbitrary number of grain species and size distributions. The size distributions of the various dust species are subdivided into different bins, resulting in a dust mixture with N_{pop} populations, each of them corresponding to a dust species and a small size bin. Assuming LTE for each individual population, the dust emissivity is given by a sum of modified blackbodies, where the temperature of each population is still determined by the condition of thermal and radiative equilibrium,

$$j_{\lambda,m}^{\text{d}} = 4\pi \rho_m V_m \sum_{i=1}^{N_{\text{pop}}} \kappa_{\lambda,i}^{\text{abs}} B_{\lambda}(T_{m,i}) \quad (21)$$

with $\kappa_{\lambda,i}^{\text{abs}}$ is the contribution of the i 'th dust population to the total absorption coefficient $\kappa_{\lambda}^{\text{abs}}$, and $T_{m,i}$ the equilibrium temperature of the i 'th population in cell number m , determined by

$$\int_0^{\infty} \kappa_{\lambda,i}^{\text{abs}} J_{\lambda} d\lambda = \int_0^{\infty} \kappa_{\lambda,i}^{\text{abs}} B_{\lambda}(T_{m,i}) d\lambda \quad (22)$$

The third option, in fact the only realistic option to model the spectral energy distribution of galaxies, is to consider NLTE dust emission. In theory, the transit from LTE to NLTE dust emission is not an enormous step. The main difference is that each dust population is not characterized by a single equilibrium temperature, but by a temperature distribution. Once the temperature distribution function has been determined, the dust emissivity can easily be determined. As argued in the Introduction, however, the practical inclusion of NLTE dust emission in radiative transfer codes is a notoriously tough nut to crack.

Rather than develop our own routines to calculate the NLTE emission for transiently heated grains, we have opted to couple the SKIRT code to the DustEM code (Compiègne et al. 2011). DustEM is a publicly available, state-of-the-art numerical tool designed to calculate the NLTE emission and extinction of dust given its size distribution, optical and thermal properties. The code builds on the work by Désert et al. (1986, 1990) and uses an adaptive temperature grid on which the temperature distribution of the grains is calculated iteratively. No LTE approximation is made, i.e. even for large grains the temperature distribution is calculated explicitly. One of the advantages why we have chosen to couple SKIRT to the DustEM code is that the latter code has been designed to deal with a variety of grain types, structures and size distributions and that new dust physics (ionization of PAHs, polarized emission, spinning dust emission, temperature-dependent dust emissivity) can easily be included. On the other hand, a consequence of choosing for a very complete and accurate NLTE routine in which basically no simplifications or assumptions have been made, is that the computation load is substantial. For a typical dust model consisting of three or four dust species each with their size distribution, the calculation of the emissivity for a single ISRF takes typically of the order of several seconds on a standard desktop/laptop computer. While this is compatible with 1D or 2D simulations with up to 10^4 cells, this is excessively long for general 3D simulations for which we have designed SKIRT.

3.5.2. A library approach for NLTE dust emission

To overcome this problem, we have adopted a strategy based on a library of dust emissivity profiles, inspired by the work of [Juvela & Padoan \(2003\)](#). Their approach consists of three steps: they first run an exploratory radiative transfer simulation on a grid with a reduced number of grid cells, without taking dust re-emission into account. This low-resolution simulation is used to determine the range of ISRFs encountered in the simulation. The second step consists of picking a small number N_{ref} of reference wavelengths λ_i^{ref} (typically $N_{\text{ref}} = 2$). The different ISRFs found in step one are discretized onto logarithmic intervals at each of the reference wavelengths, and at each bin in the N_{ref} -dimensional parameter space, the full NLTE dust emission spectrum is calculated and stored in a library. The final step in the simulation consists of running a radiative transfer simulation at the full resolution. The dust emissivity at any given cell is found by looking at the ISRF at the N_{ref} reference wavelengths and interpolating the dust emissivities from the library.

While valuable and inspiring, we see two drawbacks in the method as implemented by [Juvela & Padoan \(2003\)](#). In panchromatic radiative transfer simulations of galaxies, we typically solve the transfer equation at many UV/optical wavelengths, and hence have the ISRF at all these wavelengths at our disposal at every grid cell. It would be a pity not to use this information to determine the dust emissivity and only base our estimate on the value of the ISRF at a very small number of reference wavelengths. In particular, the ISRF in Monte Carlo simulations can be noisy in certain cells; when the dust emissivity is determined based on the value of the ISRF at a small number of reference wavelengths, this noise could lead to a significant error. Using an estimate that exploits the available information at all wavelengths can minimize this error.

The second drawback is that the library method of [Juvela & Padoan \(2003\)](#) requires an exploratory, low-resolution simulation in which the parameter space of ISRFs is explored and the library of dust emissivities is built. This extra simulation not only requires a computational overhead, it also creates the danger that it does not cover the entire range of strengths and shapes of the ISRF. For example, one can assume that the strongest ISRF in a simulation is found in small dust cells very close to the heating sources. In a low-resolution simulation, with larger dust cells, this strong ISRF will be smoothed over the larger grid cells. Similarly, the weakest ISRF (or equivalently, the coldest dust) in some simulations could be found in the inner regions of dense cores, and due to smoothing a low-resolution simulation might not reach these weakest ISRF levels. The result is that the low-resolution grid will not cover the full range of ISRFs encountered in the high-resolution grid, and hence that the library of dust emissivities must be somehow extended to incorporate this missing part in the parameter space. [Juvela & Padoan \(2003\)](#) are aware of this inconvenience (they discuss only the coldest spectrum as they concentrate on dark clouds illuminated by an external radiation field). They argue that this problem is not expected to be significant, and that it could be relieved by using a low-resolution simulation with a slightly higher density. Still, it is clear that the use of a low-resolution

grid leads to an additional overhead and complication and is a potential source of error, and it would be better to avoid it.

To overcome both problems, we have taken a slightly different approach to implement our dust emissivity library. The first step in our library approach is to calculate a number of parameters that characterize the ISRF in each dust cell after the stellar emission phase. Instead of the value of J_λ at a number of wavelengths, we use parameters that use combined information at all available wavelengths. From the various range of possibilities, we choose the lowest-order moments of $\kappa_\lambda^{\text{abs}} J_\lambda$, the product of the ISRF and the dust absorption coefficient. Instead of the actual zeroth-order moment or normalization of this function, we consider the equivalent would-be equilibrium dust temperature T_{eq} of the dust mixture, found by solving the equation

$$\int_0^\infty \kappa_\lambda^{\text{abs}} B_\lambda(T_{\text{eq}}) d\lambda = \int_0^\infty \kappa_\lambda^{\text{abs}} J_\lambda d\lambda \quad (23)$$

As a second parameter, we take the first-order moment $\kappa_\lambda^{\text{abs}} J_\lambda$, i.e. the mean wavelength,

$$\bar{\lambda} = \frac{\int_0^\infty \kappa_\lambda^{\text{abs}} J_\lambda \lambda d\lambda}{\int_0^\infty \kappa_\lambda^{\text{abs}} J_\lambda d\lambda} \quad (24)$$

In SKIRT we limit ourselves to two parameters, but in principle this procedure can be extended to more parameters.

With T_{eq} and $\bar{\lambda}$ calculated in each dust cell, we construct a 2D rectangular grid with $N_{T_{\text{eq}}} \times N_{\bar{\lambda}}$ pixels in the $(T_{\text{eq}}, \bar{\lambda})$ parameter space, based on the range of values encountered in the present simulation. In every parameter space pixel we construct a reference ISRF by averaging all ISRFs that correspond to those particular values of T_{eq} and $\bar{\lambda}$. We experimented with different ways of averaging, including taking the straight mean, the median or the mean using sigma-clipping, but found no noticeable difference. The final step of the library construction is to feed the reference ISRFs to the DustEM routine, and save each of the resulting dust emissivity profiles in the library.

Once the library is constructed, finding the correct dust emissivity for a given dust cell is straightforward, as each dust cell is already connected to a certain pixel in the $(T_{\text{eq}}, \bar{\lambda})$ parameter space and hence a dust emissivity profile in the library.

4. IMPLEMENTATION DETAILS

While the very first version of SKIRT were written in Fortran 77, the code is now completely written in ANSI C++ and currently contains some 30 000 lines of code. It uses the object-oriented nature of the C++ language extensively to support a strong modularity. The use of inheritance and abstract classes renders the inclusion of new components (such as new density distributions for the stars or dust, or new dust mixtures) straightforward. The DustEM code is written in Fortran 95 and has been slightly adapted to be coupled to SKIRT. The entire SKIRT code is driven by a graphic user interface written in PyQt. Batch jobs can be run using a command line version with XML input files.

An important implementation aspect of SKIRT is the parallelism. Parallelism can typically work on two domains: data parallelism focuses on distributing the data

across different parallel computing nodes, with the principle aim of enabling simulations that need more memory consumption than is available on a single node. Task or control parallelism focuses on distributing execution processes (threads) across different parallel computing nodes with the principle aim of decreasing the run-time of a program. Ideally, both approaches can be combined. Monte Carlo radiative transfer codes are easily parallelized in a task parallelism approach: the different levels of iterations can easily be split over different nodes. SKIRT uses the OpenMP protocol to support task parallelism on shared-memory machines. One of the advantages of OpenMP parallelism is the spectacularly low coding cost: less than 100 lines of code (on a total of 30 000) have been added to SKIRT to convert it from a serial into a parallel code. The main parallelism is situated in the loop over wavelength, which implies that SKIRT runs both the stellar and dust phases at different wavelengths simultaneously.

One of the planned future developments of SKIRT is to look into the possibilities of using the MPI interface for data parallelism, in order to allow Monte Carlo simulations to be run on distributed memory systems. Data parallelism is much harder to achieve for Monte Carlo simulations than task parallelism. The main reason is the non-local nature of the physical problem: each photon package in a simulation typically requires data from the entire physical domain (read access to calculate the optical depth and write access to update the absorption rates). Contrary to e.g. grid-based hydrodynamic codes, distributing the dust cells over different parts of memory would imply an enormous overhead in communication between the different nodes. In principle, data parallelism could be achieved by splitting the data in the wavelength dimension, where different nodes contain different parts of the absorption rate counters of the dust grid and different parts of the detectors. However, in this approach a large amount of data (such as the dust density grid) would need to be shared/copied between the different nodes and communication overheads would be significant in the dust emission phase. Future work will investigate whether the benefits of distributed-memory parallelism can outweigh the communication overheads and the significant additional coding complexity.

5. TESTS AND APPLICATIONS

We have run extensive tests to check the accuracy of the SKIRT code against other radiative transfer codes. The early versions of the code were already tested against several other results, most importantly the set of spiral galaxy models by [Byun et al. \(1994\)](#) (see e.g. [Baes et al. 2003](#)). We have successfully tested the LTE version of SKIRT against the 1D and 2D LTE circumstellar benchmark problems of [Ivezic et al. \(1997\)](#) and [Pascucci et al. \(2004\)](#). SKIRT is also one of the codes used in a new ongoing LTE benchmark effort focusing on a disc galaxy environment ([Baes et al. 2011](#), in preparation). The preliminary results, based on the results from five independent radiative transfer codes indicate excellent agreement, with relative differences in the SEDs around the 1% level or even below.

As a full NLTE radiative transfer benchmark is not (yet) available at the moment, we have tested our NLTE radiative transfer code, and particularly the library ap-

proach, using different models with gradually increasing levels of complexity. In order to run simulations in a realistic setting, we adopt the Sbc galaxy UGC 4754 as a template model. UGC 4754 is a edge-on spiral galaxy, which has always been a favorite class of galaxies for radiative transfer modellers, as the dust is clearly visible both in absorption and emission (e.g. [Xilouris et al. 1997, 1998, 1999](#); [Popescu et al. 2000, 2011](#); [Misiriotis et al. 2001](#); [Dasyra et al. 2005](#); [Bianchi 2007, 2008](#)). This galaxy was one of the first large edge-on galaxies to be observed with the Herschel Space Observatory as part of the Herschel Astrophysical Terahertz Large Area Survey (H-ATLAS, [Eales et al. 2010](#)) science demonstration phase observations. In [Baes et al. \(2010\)](#), we fitted a radiative model to the observed images of UGC 4754 in the SDSS and UKIDSS bands. We subsequently used the SKIRT code to predict the galaxy’s SED and images at FIR wavelengths. While the radiative transfer model used in that paper was sufficient to serve its goal (investigation of the dust energy balance), it suffered from two significant limitations: the assumptions of LTE dust emission and of a smooth interstellar medium. These two assumptions prevented us from making a self-consistent model covering the entire spectral region from the UV to mm wavelengths. Moreover, they also might introduce a significant source of uncertainty, as it has been demonstrated by several authors that the extinction properties of a clumpy interstellar medium can be significantly different from those of a smooth medium (e.g. [Hobson & Scheuer 1993](#); [Witt & Gordon 1996, 2000](#); [Wolf et al. 1998](#); [Bianchi et al. 2000](#); [Matthews & Wood 2001](#); [Pierini et al. 2004](#); [Doty et al. 2005](#)).

In this section we will gradually refine our model for UGC 4754 from a smooth, 2D, LTE model to a fully 3D model that includes NLTE dust emission and a clumpy structure of the dusty ISM. The main objectives are to test the accuracy of our approach using a realistic setting, and to demonstrate the ability of SKIRT to run realistic 3D NLTE radiative transfer calculations. For a full investigation of the dust energy balance in spiral galaxies, based on our refined SKIRT code and multi-wavelength imaging data, we refer to future work.

5.1. 2D radiative transfer models

The starting point for our models is the best fitting, smooth 2D model from [Baes et al. \(2010\)](#). The stellar distribution consists of two components: a double-exponential stellar disc with a scale length of 4.05 kpc and a scale height of 330 pc, and a flattened Sérsic bulge with a major axis effective radius of 800 pc, a Sérsic parameter of 0.9 and an intrinsic flattening of 0.6. Both components have a similar intrinsic SED, corresponding to a population of 8 Gyr old with an exponentially decaying star formation rate and an initial burst duration of 0.15 Gyr. The total bolometric luminosity of the system is $1.8 \times 10^{10} L_{\odot}$, of which the bulge contributes 8%. The dust is also distributed in a double-exponential disc with a scale length of 6.1 kpc and a scale height of 270 pc. The total amount of dust is characterized by the *g*-band edge-on optical depth of 0.73, which corresponds to a total dust mass of $1.0 \times 10^7 M_{\odot}$. Contrary to [Baes et al. \(2010\)](#) where we used the [Zubko et al. \(2004\)](#) model to describe the dust optical properties, we now use the [Draine & Li \(2007\)](#) dust model, as this model is embed-




FIG. 2.— SEDs of the SKIRT 2D models for UGC 4754. The blue curve corresponds to a model assuming LTE dust re-emission, the black curve shows the SED of the model that includes NLTE dust re-emission. For the latter model, the yellow line shows the contribution of the thermal dust emission to the SED. The red curve represents the intrinsic flux of the model without dust extinction, and the data points are the observed total fluxes of UGC 4754.

ded in the DustEM library (Compiegne et al. 2011).

Simulations are run on a wavelength grid with 181 grid points, with 101 grid points distributed logarithmically between 0.05 and 5000 μm and 80 additional grid points distributed logarithmically between 1 and 30 μm to capture the PAH peaks. For our 2D simulations, we considered an axisymmetric grid with 51 grid points in the radial direction and 51 grid points in the vertical direction, resulting in a total number of $N_{\text{cells}} = 2500$ grid cells. The grid points are chosen to have a power-law distribution, with an extent of 30 kpc (2 kpc) for the radial (vertical) distribution and a size ratio of 30 between the innermost and outermost bins. In all SKIRT runs discussed here, we used 10^6 photon packages for each wavelength in both the stellar and the dust emission phase.

Figure 2 shows the resulting SEDs of two different SKIRT 2D simulations based on this model setup, as well as the observed GALEX, SDSS, UKIDSS, IRAS and Herschel fluxes for UGC 4754. The blue curve shows the SED corresponding to a model assuming LTE dust re-emission, where we took into account that different grain types and sizes reach different equilibrium temperatures. The black curve shows the SED of the model that includes NLTE dust re-emission using the library approach discussed in Section 3.5. The red curve represents the flux of the model without dust extinction, the yellow line corresponds to the contribution of the dust to the SED in the NLTE case. It is logical that there is no difference between the LTE and NLTE models in the UV/optical/NIR part of the SED, as only the dust emissivity differs between the models. There is, however, a significant difference between the LTE and NLTE models in the MIR/FIR/submm window: in the LTE models, all absorbed radiation is re-emitted as modified blackbody emission in the FIR/submm region, whereas the NLTE




FIG. 3.— The relative difference $(F^{\text{bf}} - F^{\text{lib}})/F^{\text{lib}}$ between the dust SEDs of 2D NLTE models for UGC 4754 using the brute force (bf) approach and the library (lib) approach.

models emit part of the absorbed radiation in the MIR region.

To test the accuracy of our library approach, we ran a second NLTE simulation, where we did not use the library approach, but where we calculated the dust emissivity in each cell by an explicit call to the DustEM code. The relative difference between the SEDs of the SKIRT NLTE models using the library approach and the brute force approach is shown in Figure 3. This figure convincingly shows that the library approach is accurate: the relative error is everywhere below 0.2% and in the region where the dust emission dominates even smaller. The difference in run time between the two approaches is substantial, and this difference is only due to the difference in the number of calls to the DustEM routine. In the brute force approach, the number of calls is obviously equal to N_{cells} (or sometimes a bit less if there are empty dust cells). The maximum number of calls to the DustEM routine in the library approach is obviously $N_{T_{\text{eq}}} \times N_{\lambda}$, the number of pixels in the library parameter space. As the intensity of the ISRF is the main parameter that drives the shape of the dust emissivity spectrum (Draine & Li 2001; Compiegne et al. 2011),




FIG. 4.— Comparison of our library approach and the Milky Way ISRF template approach for 2D NLTE models for UGC 4754. The top panel shows the total SEDs and the dust emission SEDs corresponding to both approaches, the top panel shows the relative difference $(F_\lambda^{\text{tem}} - F_\lambda^{\text{lib}})/F_\lambda^{\text{lib}}$ between the dust SEDs using the MW ISRF template (tem) approach and our library (lib) approach.

it is appropriate to choose $N_{T_{\text{eq}}}$ somewhat larger than $N_{\bar{\lambda}}$; the values we propose are $N_{T_{\text{eq}}} = 25$ and $N_{\bar{\lambda}} = 10$. This infers a maximum of 250 calls to the DustEM routine. In most cases, however, not the entire $(T_{\text{eq}}, \bar{\lambda})$ parameter space is covered, such that the number of calls is even further reduced. In the present simulation, we needed to call the DustEM routine only 71 times in the library approach, compared to 2500 times in the brute force approach. Given that the average DustEM run time is about 7 s (on a typical desktop computer) and that the overhead of the construction of the library is negligible, this makes a substantial difference. It should be noted that 2500 calls is still a very manageable number, but the brute force approach becomes impossible when moving to 3D grids with several million cells. For example, for a simulation with ten million dust cells, we would need a DustEM computation time of more than 2 years on a single core computer, compared to several minutes using the library approach.

As a final test and a demonstration of the necessity of our library approach, we ran our 2D simulation again, but now using a set of precomputed template emissivity profiles. In a sense this mimics the approach taken by e.g. Wood et al. (2008), although their approach is somewhat different as they calculate the emission by big grains using the LTE approximation and only use a template for the small grains. The basic idea of our template approach is that the dust emissivity from a cell depends only on a single parameter U , being the integrated mean intensity of the ISRF in the considered cell expressed in terms of the integrated mean intensity of the ISRF in the Milky Way,

$$U = \frac{\int_0^\infty J_\lambda d\lambda}{\int_0^\infty J_\lambda^{\text{MW}} d\lambda} \quad (25)$$

For the ISRF of the MW, the standard parameterization of Mathis et al. (1983) is adopted. We constructed

a library of 501 dust emissivity profiles, corresponding to values of U distributed logarithmically between $U_{\text{min}} = 10^{-5}$ and $U_{\text{max}} = 10^5$. These dust emissivity profiles can be computed once and for all (using the DustEM routine) and saved in a file. In the dust emission phase, we simply calculate U in every dust cell and determine the dust emissivity profile using logarithmic interpolation between the precomputed library profiles. This approach definitely has a strong appeal: it is straightforward to implement and very fast, as it requires only the calculation of U and a simple interpolation of precomputed values. In this sense it is more attractive than our library approach, which requires the construction of a library for every simulation. The disadvantage of the template approach, however, is that only the strength and not the shape or hardness of the ISRF is taken into account to calculate the dust emissivity. As also argued by Jonsson et al. (2010), the shape of the ISRF can have a significant importance, both because of the differing cross-sections of grains of different sizes and because high-energy photons will excite larger thermal fluctuations than low-energy photons for a given value of U .

In Figure 4 we show the comparison of the fixed template approach and our dynamic library approach (dynamic in the sense that the library is tailored to the specific simulation). The top panel shows the total SED and the dust emission SED for our 2D model for UGC 4754 obtained using both approaches. At first sight, the SEDs agree fairly well (definitely when plotted in log-log scale). Looking at the bottom panel, where we plot the relative difference of the dust emission SED corresponding to both approaches, we do see a significant difference, with relative deviations up to more than 40%. Here we have to give an important side note, in the sense that we believe that this 40% difference is actually a underestimate of the error one can make. The shape of the intrinsic SED of the stellar population in our model is independent of position; as a result, the variations in the shape of the ISRF in different cells in the simulation are only the result of varying levels of absorption and scattering. Since the stellar population model we adopted for UGC 4754 (an 8 Gyr old population with an exponentially decaying star formation rate and an initial burst duration of 150 Myr) is not very different from the average stellar population in the Milky Way, this implies that the shape of the ISRF in the different cells in our model will on average be quite close to the shape of the Milky Way ISRF. As far as the comparison between our library approach and the Milky Way template approach concerns, we must hence conclude that UGC 4754 simulations are not the strongest test. For systems with ISRFs which deviate much more from the average Milky Way ISRF, such as starburst galaxies or circumstellar discs around young hot stars, we expect much larger differences than the 40% we obtained here.

5.2. 3D clumpy radiative transfer models

Having tested the accuracy of our library approach, we are ready to run full-scale 3D NLTE radiative transfer models with SKIRT. The first 3D model we consider has exactly the same set-up as the 2D NLTE model, except that we now consider a uniform 3D cartesian grid. In the x and y directions we consider 401 grid points each and a maximum extent of 30 kpc, in the vertical direction




FIG. 5.— The relative difference $(F_{\lambda}^{3D} - F_{\lambda}^{2D})/F_{\lambda}^{2D}$ between the SEDs of 2D and 3D NLTE models for UGC 4754. The red curve corresponds to a 3D model with a uniform cartesian dust grid in all three dimensions, the blue curve corresponds to a 3D model with a power-law grid in the vertical direction, similar to the vertical grid structure of the 2D model.

we use 61 grid points with a maximum extent of 2 kpc. This results in $N_{\text{cells}} = 9.6 \times 10^6$ grid cells, each with a dimension of 150 pc in the x and y directions and 66.7 pc in the vertical direction. Note that we need to store, at each dust grid cell, the entire ISRF J_{λ} at each of the wavelength grid points, which basically turns our grid into a 4-dimensional grid structure with 1.73×10^9 grid cells. The memory required to run such a large-scale SKIRT radiative transfer simulation is about 23 GB.

Figure 5 compares the SED of the smooth 2D and 3D models for UGC 4754. The relative differences are below 2% in the entire UV-mm domain. The existing small differences are mainly due to the discretization of the grid in the vertical direction: for the 2D model we used a vertical grid with a power-law distribution, with smaller bins close to the equatorial plane, where the dust density has the strongest gradients. The innermost grid cell has a height of only 8.75 pc, compared to the 66.7 pc in the case of the 3D grid. The result is that the discretization of the dust density on the 3D grid is much coarser. To demonstrate that this vertical grid distribution is the origin of this $< 2\%$ difference, we ran another 3D simulation where we now apply the same power-law distribution for the vertical grid cells as we did for the 2D model. The relative differences between the SED of this model and the SED of the 2D simulation are also indicated in Figure 5.

The reason why we considered a uniform cartesian dust grid is because such a grid forms the basis for a fully 3D model with a clumpy, two-phase dust distribution. To generate such a model, we followed the strategy outlined by Witt & Gordon (1996). The dusty interstellar medium consists of two phases, a smooth inter-clump component and a clumpy component, and is characterized in terms of two parameters, namely the volume filling factor f_f of the dense clumps and the density contrast C between the clump and inter-clump medium. The practical construction of the two-phase medium consists of randomly assigning a status (clump or inter-clump) to each dust cell in the dust medium. Typical values for the parameters C and f_f vary widely in the literature. We use the values $C = 100$ and $f_f = 0.1$ in this work, which are within the range of typical parameters used in other studies (e.g. Kuchinski et al. 1998; Witt & Gordon 2000; Bianchi et al. 2000; Matthews & Wood 2001; Pierini et al. 2004; Doty et al. 2005).

In Figure 6, we show simulated images corresponding to the clumpy and smooth 3D models for the inner region

of UGC 4754 at three different wavelengths (0.55, 24 and 500 μm). It is important to note that the pixel-to-pixel variations in the images on the bottom row are not due to Poisson noise in the Monte Carlo routine, but represent real intensity variations due to the clumpy nature of the dust in the models. Figure 7 compares the total SED of the clumpy and smooth models. The relative differences are below 6% in the entire UV-mm domain. At UV wavelengths, the clumpy model is slightly brighter than the smooth model, or put differently, slightly less UV radiation is absorbed by the dust. This is in agreement with conclusions found by other authors: for a fixed amount of dust, a clumpy dust medium absorbs radiation less efficiently than a smooth dust medium (e.g. Bianchi et al. 2000; Pierini et al. 2004). When moving to longer wavelengths, the difference between the smooth and clumpy models decreases, as the fraction of absorbed versus unattenuated radiation decreases with increasing wavelength. At NIR wavelengths, there is virtually no difference anymore between the SEDs of the smooth and clumpy models. Moving to MIR and FIR wavelengths, we find that the clumpy model emits significantly less, particularly at wavelengths up to about 100 μm . This is no surprise, as the clumpy model was less efficient in absorbing UV and optical radiation. The result is that, for the same total mass, the dust in a clumpy model is on average both cooler and less luminous than the dust in a smooth model.

6. CONCLUSIONS

We have presented an updated version of the 3D Monte Carlo radiative transfer code SKIRT. The code uses various advanced optimization techniques, both well-known and novel ones, that make the Monte Carlo process orders of magnitude more efficient than the most basic Monte Carlo technique. These techniques include an optimized combination of eternal forced scattering and continuous absorption, a multi-gaussian expansion technique and an efficient foam generator to generate random positions from the stellar density, and the use of peeling-off and smart detectors to create high signal-to-noise images and SEDs.

The main novelty of the new SKIRT code is the possibility to calculate the dust temperature distribution and the associated infrared and submm emission with a full incorporation of the emission of transiently heated grains and PAH molecules. To achieve this, we have chosen to link the SKIRT code to DustEM (Compiègne et al. 2011), a publicly available, state-of-the-art numerical tool designed to calculate the NLTE emission of arbitrary mixtures of dust grains. The advantages of this approach is that no LTE approximation is made, even for large grains, and that new physics (such as spinning dust emission or a temperature-dependent dust emissivity) can readily be included. We have implemented a library approach to limit the computational cost of the NLTE dust emission calculations inherent in DustEM. Our approach is inspired by the work by Juvela & Padoan (2003), but uses a slightly different approach that makes maximum use of all information in the simulation to calculate the dust emissivity and avoids the need for additional low-resolution simulations.

We have tested the accuracy of the SKIRT code, in particular of our NLTE library approach, through a set




FIG. 6.— Simulated images corresponding to the smooth (top row) and clumpy (bottom row) 3D models for UGC 4754. The three panels correspond to three different regimes: stellar emission and dust extinction ($0.55 \mu\text{m}$, left), NLTE emission by small grains ($24 \mu\text{m}$, middle) and LTE emission by big grains ($500 \mu\text{m}$, right).




FIG. 7.— The relative difference $(F_{\lambda}^{\text{clumpy}} - F_{\lambda}^{\text{smooth}})/F_{\lambda}^{\text{smooth}}$ between the SEDs of clumpy and smooth 3D NLTE models for UGC 4754.

of simulations based on the edge-on spiral galaxy UGC 4754, previously modelled by Baes et al. (2010). The models we ran were gradually refined from a smooth, 2D, LTE model to a fully 3D model that includes NLTE dust emission and a clumpy structure of the dusty ISM.

Using 2D models, we demonstrated the accuracy of our library approach: the relative differences in the SED between a model that uses the library approach and a model that uses brute force to calculate the dust emission are less than 0.2% at all wavelengths. Even for this 2D model with only 2500 dust cells, the difference in run time between both approaches are substantial; for 3D grids with several million dust cells the brute force approach becomes impossible. We have also explored the possibility to use a fixed set of precomputed dust emission templates instead of a dynamic library as the one we have chosen. While a template approach has the advantage that it is easier to implement and faster to run, we have demonstrated that it leads to significant deviations due to the fact that it does not take into account the shape of the interstellar radiation field. This highlights the need for a more advanced approach such as the

library approach we propose.

We have subsequently applied the SKIRT code to calculate full-scale 3D NLTE models for UGC 4754. We found small differences ($< 2\%$) between 2D and 3D smooth models that are mainly due to differences in the vertical discretization of the internal grid. Finally, we have compared 3D models with a smooth and a clumpy interstellar dust medium. We confirm the result found by other authors that, for a fixed amount of interstellar dust, a clumpy dust medium absorbs radiation less efficiently than a smooth dust medium. As a direct consequence, the dust in clumpy models is on average both cooler and less luminous, and the observed infrared emission of clumpy models is less than the emission at these wavelengths of smooth models with the same dust mass.

Our simulations demonstrate that, given the appropriate use of optimization techniques, it is possible to efficiently and accurately perform Monte Carlo radiative transfer simulations of arbitrary 3D structures of several million dust cells, including a full calculation of the NLTE emission by arbitrary dust mixtures. This significantly increases the number of applications where detailed radiative transfer modeling can be used. For example, we have started an investigation of the energy balance crisis in a set of edge-on spiral galaxies: our intention is to fit detailed radiative transfer models to UV/optical/NIR images for a set of edge-on spiral galaxies, predict the resulting MIR/FIR/submm emission and compare these predictions with the available long wavelength data. Many other applications (AGNs, circumstellar discs, merging galaxies,...) are possible, and the authors welcome all projects that can make use of SKIRT.

M. B., J. V. and E. V. P. acknowledge the support

of the Fund for Scientific Research Flanders (FWO-Vlaanderen) that made the development of the SKIRT code possible. M. B., I. D. L. and J. F. thank the Belgian Federal Science Policy Office (BELSPO) for the financial support. M. S. acknowledges support of the Ministry of Education and Science of the Republic of Serbia through the projects “Astrophysical Spectroscopy of Extragalac-

tic Objects” (176001) and “Gravitation and the Large Scale Structure of the Universe” (176003) and support of the European Commission (Erasmus Mundus Action 2 partnership between the European Union and the Western Balkans, <http://www.basileus.ugent.be>) during his mobility period at Ghent University.

REFERENCES

- Andonov, A., Arbuzov, A. B., Bondarenko, S. G., Christova, P., Kolesnikov, V. A., Nanava, G., Sadykov, R. R. 2010, *Physics of Atomic Nuclei*, 73, 1761
- Baes, M. 2008, *MNRAS*, 391, 617
- Baes, M., & Dejonghe, H. 2001, *MNRAS*, 326, 733
- Baes, M., & Dejonghe, H. 2001, *ApJ*, 563, L19
- Baes, M., & Dejonghe, H. 2002, *MNRAS*, 335, 441
- Baes, M., et al. 2003, *MNRAS*, 343, 1081
- Baes, M., Dejonghe, H., & Davies, J. I. 2005a, *The Spectral Energy Distributions of Gas-Rich Galaxies: Confronting Models with Data*, AIP Conference Series, 761, 27
- Baes, M., Stamatellos, D., Davies, J. I., Whitworth, A. P., Sabatini, S., Roberts, S., Linder, S. M., & Evans, R. 2005b, *NewA*, 10, 523
- Baes, M., et al. 2010, *A&A*, 518, L39
- Bendo, G. J., et al. 2010, *A&A*, 518, L65
- Bianchi, S. 2007, *A&A*, 471, 765
- Bianchi, S. 2008, *A&A*, 490, 461
- Bianchi, S., Ferrara, A., & Giovanardi, C. 1996, *ApJ*, 465, 127
- Bianchi, S., Ferrara, A., Davies, J. I., & Alton, P. B. 2000, *MNRAS*, 311, 601
- Bruzual A., G., Magris, G., & Calvet, N. 1988, *ApJ*, 333, 673
- Byun, Y. I., Freeman, K. C., & Kylafis, N. D. 1994, *ApJ*, 432, 114
- Cappellari, M. 2002, *MNRAS*, 333, 400
- Cappellari, M., et al. 2006, *MNRAS*, 366, 1126
- Cashwell, E. D., & Everett, C. J., 1959, *A Practical Manual on the Monte Carlo Method for Random Walk Problems*, Pergamon, New York
- Clements, D. L., et al. 2010, *A&A*, 518, L8
- Compiègne, M., et al. 2011, *A&A*, 525, A103
- Dasyra K. M., Xilouris E. M., Misiriotis A., Kylafis N. D., 2005, *A&A*, 437, 447
- De Looze, I., et al. 2010, *A&A*, 518, L54
- De Looze, I., Baes, M., Fritz, J., Verstappen, J., 2011, *MNRAS*, submitted
- Désert, F. X., Boulanger, F., & Shore, S. N. 1986, *A&A*, 160, 295
- Désert, F.-X., Boulanger, F., & Puget, J. L. 1990, *A&A*, 237, 215
- de Vaucouleurs, G. 1948, *Annales d’Astrophysique*, 11, 247
- di Bartolomeo, A., Barbaro, G., & Perinotto, M. 1995, *MNRAS*, 277, 1279
- Disney, M., Davies, J., & Philipps, S. 1989, *MNRAS*, 239, 939
- Doty, S. D., Metzler, R. A., & Palotti, M. L. 2005, *MNRAS*, 362, 737
- Draine, B. T., & Li, A. 2001, *ApJ*, 551, 807
- Draine, B. T., & Li, A. 2007, *ApJ*, 657, 810
- Draine, B. T., et al. 2007, *ApJ*, 663, 866
- Dwek, E. 1986, *ApJ*, 302, 363
- Eales, S., et al. 2010, *PASP*, 122, 499
- Emsellem, E., Monnet, G., & Bacon, R. 1994a, *A&A*, 285, 723
- Emsellem, E., Monnet, G., Bacon, R., & Nieto, J.-L. 1994b, *A&A*, 285, 739
- Emsellem, E. 1995, *A&A*, 303, 673
- Emsellem, E., Dejonghe, H., & Bacon, R. 1999, *MNRAS*, 303, 495
- Fischer, O., Henning, T., & Yorke, H. W. 1994, *A&A*, 284, 187
- Galametz, M., et al. 2010, *A&A*, 518, L55
- Gomez, H. L., et al. 2010, *A&A*, 518, L45
- Gordon, K. D., Misselt, K. A., Witt, A. N., & Clayton, G. C. 2001, *ApJ*, 551, 269
- Guhathakurta, P., & Draine, B. T. 1989, *ApJ*, 345, 230
- Haas, T., Makarenko, V. 2011, *The European Physical Journal C*, 7, 1574
- Helou, G., Lu, N. Y., Werner, M. W., Malhotra, S., & Silberman, N. 2000, *ApJ*, 532, L21
- Hernquist, L. 1990, *ApJ*, 356, 359
- Hobson, M. P., & Scheuer, P. A. G. 1993, *MNRAS*, 264, 145
- Inoue, A. K. 2005, *MNRAS*, 359, 171
- Ivezic, Z., Groenewegen, M. A. T., Men’shchikov, A., & Szczerba, R. 1997, *MNRAS*, 291, 121
- Jadach, S. 2003, *Computer Physics Communications*, 152, 55
- Jadach, S., Sawicki, P. 2007, *Computer Physics Communications*, 177, 441
- Jadach, S., Skrzypek, M. 2006, *Computer Physics Communications*, 175, 511
- Jaffe, W. 1983, *MNRAS*, 202, 995
- Jonsson, P. 2006, *MNRAS*, 372, 2
- Jonsson, P., Groves, B. A., & Cox, T. J. 2010, *MNRAS*, 403, 17
- Juvela, M. 2005, *A&A*, 440, 531
- Juvela, M., & Padoan, P. 2003, *A&A*, 397, 201
- Kennicutt, R. C., Jr., et al. 2003, *PASP*, 115, 928
- Kuchinski, L. E., Terndrup, D. M., Gordon, K. D., & Witt, A. N. 1998, *AJ*, 115, 1438
- Kylafis N. D., Bahcall J. N., 1987, *ApJ*, 317, 637
- Lucy, L. B. 1999, *A&A*, 344, 282
- Mathis, J. S., Mezger, P. G., & Panagia, N. 1983, *A&A*, 128, 212
- Matthews, L. D., & Wood, K. 2001, *ApJ*, 548, 150
- Mattila, K. 1970, *A&A*, 9, 53
- Misselt, K. A., Gordon, K. D., Clayton, G. C., & Wolff, M. J. 2001, *ApJ*, 551, 277
- Misiriotis, A., Popescu, C. C., Tuffs, R., & Kylafis, N. D. 2001, *A&A*, 372, 775
- Monnet, G., Bacon, R., & Emsellem, E. 1992, *A&A*, 253, 366
- Niccolini, G., Woitke, P., & Lopez, B. 2003, *A&A*, 399, 703
- Oliver, S. J., et al. 2010, *A&A*, 518, L21
- Pascucci, I., Wolf, S., Steinacker, J., Dullemond, C. P., Henning, T., Niccolini, G., Woitke, P., & Lopez, B. 2004, *A&A*, 417, 793
- Pierini, D., Gordon, K. D., Witt, A. N., & Madsen, G. J. 2004, *ApJ*, 617, 1022
- Pinte, C., Ménard, F., Duchêne, G., & Bastien, P. 2006, *A&A*, 459, 797
- Plummer, H. C. 1911, *MNRAS*, 71, 460
- Pohlen, M., et al. 2010, *A&A*, 518, L72
- Popescu, C. C., Misiriotis, A., Kylafis, N. D., Tuffs, R. J., & Fischera, J. 2000, *A&A*, 362, 138
- Popescu, C. C., Tuffs, R. J., Dopita, M. A., Fischera, J., Kylafis, N. D., & Madore, B. F. 2011, *A&A*, 527, A109
- Popović, L. C., Jovanović, P., Stalevski, M., Antoniu, S., Andrei, A. H., Kovačević, J., Baes, M. 2011, *A&A*, submitted
- Rigby, E. E., et al. 2011, *MNRAS*, 955
- Roussel, H., et al. 2010, *A&A*, 518, L66
- Semionov, D., & Vansevicius, V. 2005, *Baltic Astronomy*, 14, 543
- Sérsic, J. L. 1968, *Cordoba, Argentina: Observatorio Astronomico Siebenmorgen*, R., Kruegel, E., & Mathis, J. S. 1992, *A&A*, 266, 501
- Smith, J. D. T., et al. 2007, *ApJ*, 656, 770
- Smith, M. W. L., et al. 2010, *A&A*, 518, L51
- Stalevski, M., Fritz, J., Baes, M., Nakos, Th., Popović, L. Č. 2011, *MNRAS*, submitted
- Stamatellos, D., & Whitworth, A. P. 2003, *A&A*, 407, 941
- Steinacker, J., Henning, T., Bacmann, A., & Semenov, D. 2003, *A&A*, 401, 405
- van den Bosch, R. C. E., van de Ven, G., Verolme, E. K., Cappellari, M., & de Zeeuw, P. T. 2008, *MNRAS*, 385, 647
- Vidal, E., & Baes, M. 2007, *Baltic Astronomy*, 16, 101
- Witt, A. N. 1977, *ApJS*, 35, 1
- Witt, A. N., Thronson, H. A., Jr., & Capuano, J. M., Jr. 1992, *ApJ*, 393, 611
- Witt, A. N., & Gordon, K. D. 1996, *ApJ*, 463, 681
- Witt, A. N., & Gordon, K. D. 2000, *ApJ*, 528, 799
- Wolf, S. 2003, *Computer Physics Communications*, 150, 99
- Wolf, S., Fischer, O., & Pfau, W. 1998, *A&A*, 340, 103
- Wood, K., Whitney, B. A., Robitaille, T., & Draine, B. T. 2008, *ApJ*, 688, 1118
- Xilouris, E. M., et al. 1997, *A&A*, 325, 135
- Xilouris, E. M., et al. 1998, *A&A*, 331, 894

Xilouris, E. M., Byun, Y. I., Kylafis, N. D., Paleologou, E. V., &
Papamastorakis, J. 1999, *A&A*, 344, 868
Yusef-Zadeh, F., Morris, M., & White, R. L. 1984, *ApJ*, 278, 186

Zubko, V., Dwek, E., & Arendt, R. G. 2004, *ApJS*, 152, 211